

WHAT PRACTITIONERS NEED TO KNOW . . .

. . . About Event Studies

Mark P. Kritzman

Event studies measure the relationship between an event that affects securities and the return of those securities. Some events, such as a regulatory change or an economic shock, affect many securities contemporaneously; other events, such as a change in dividend policy or a stock split, are specific to individual securities.

Event studies are often used to test the efficient market hypothesis. For example, abnormal returns that persist after an event occurs or abnormal returns that are associated with an *anticipated* event contradict the efficient market hypothesis. Aside from tests of market efficiency, event studies are valuable in gauging the magnitude of an event's impact.

A classic event study published in 1969 by Fama, Fisher, Jensen, and Roll examined the impact of stock splits on security prices.¹ The authors found that abnormal returns dissipated rapidly following the news of stock splits, thus lending support to the efficient market hypothesis.

How to Perform An Event Study in Seven Easy Steps

The following steps describe one of several approaches for conducting an event study of a firm-specific event:

■ *Define the event and identify the timing of its occurrence.* The timing of the event is not necessarily the period during which the event occurs. Rather, it may be the investment period immediately preceding the announcement of the event.

■ *Arrange the security performance data relative to the timing of the event.* If information about the event is released fully on a specific day with time remaining for traders to react, the day of the announcement is period zero. Then, measurement periods preceding and following the event are selected. For example, if the 90 trading days preceding the event and the 10 days following the event are designated as the pre- and post-event periods, the pre-event trading days would be la-

beled $t - 90, t - 89, t - 88, \dots, t - 1$; the event day, $t = 0$; and the post-event trading days, $t + 1, t + 2, t + 3, \dots, t + 10$. Because the event is specific to each security, these days will differ across securities in calendar time.

■ *Separate the security-specific component of return from the security's total return during the pre-event measurement period.* One approach is to use the market model to isolate security-specific return. First, each security's daily returns during the pre-event measurement period from $t - 90$ through $t - 1$ are regressed on the market's returns during the same period. The security-specific returns are defined as the differences between the security's daily returns and the daily returns predicted from the regression equation (the security's alpha plus its beta times the market's daily returns). This calculation is described by Equation 1:

$$A_{i,t} = R_{i,t} - \hat{\alpha}_i - \hat{\beta}_i(R_{m,t}), \quad (1)$$

where

$A_{i,t}$ = security-specific return of security i in period t

$R_{i,t}$ = total return of security i in period t

$\hat{\alpha}_i$ = alpha of security i estimated from pre-event measurement period

$\hat{\beta}_i$ = beta of security i estimated from pre-event measurement period

$R_{m,t}$ = total return of market in period t

■ *Estimate the standard deviation of the daily security-specific returns during the pre-event measurement period from $t - 90$ through $t - 1$.* This calculation is shown in Equation 2:

$$\sigma_{i,pre} = \sqrt{\frac{\sum_{t=-90}^{-1} (A_{i,t} - A_{i,pre})^2}{n - 1}}, \quad (2)$$

where

$\sigma_{i,pre}$ = standard deviation of security-specific returns of security i estimated from pre-event measurement period

Mark P. Kritzman, CFA, is a Partner of Windham Capital Management in Cambridge, Massachusetts.

$A_{i,pre}$ = average of security-specific returns of security i estimated from pre-event measurement period
 n = number of days in pre-event measurement period

■ *Isolate the security-specific return during the event and post-event periods.* To estimate the security-specific return each day during these periods, subtract from each security's total return each day the security's alpha and beta times the market's return on that day. The alphas and betas are the same as those estimated from the pre-event regressions. The equation for estimating these returns is the same as Equation 1. The subscript t , however, ranges from 0 to +10 rather than from -90 to -1.

■ *Aggregate the security-specific returns and standard deviations across the sample of securities on the event day and the post-event days;* that is, sum the security-specific returns for each day and divide by the number of securities in the sample, as shown in Equation 3:

$$\bar{A}_t = \frac{\sum_{i=1}^N A_{i,t}}{N}, \quad (3)$$

where

\bar{A}_t = average across all securities of security-specific returns in period t
 N = number of securities in sample

The standard deviations are aggregated by squaring the standard deviation of each security's specific return estimated during the pre-event period, summing these values across all securities, taking the square root of this sum, and then dividing by the number of securities. Equation 4 shows this calculation:

$$\sigma_{N,pre} = \sqrt{\frac{\sum_{i=1}^N \sigma_{i,pre}^2}{N}}, \quad (4)$$

where

$\sigma_{N,pre}$ = aggregate of pre-event standard deviations of security-specific returns across all securities

■ *Test the hypothesis that the security-specific returns on the event day and post-event days differ significantly from zero.* The t -statistic is computed by dividing the average of the security-specific returns across all securities each day by the aggrega-

tion of the standard deviations across all securities as described in the previous step. Then, depending on the degrees of freedom, determine whether the event significantly affects returns. That is,

$$t\text{-statistic} = \frac{\bar{A}_t}{\sigma_{N,pre}}. \quad (5)$$

If the event is unanticipated and the t -statistic is significant on the day of the event but insignificant on the days following the event, a reasonable conclusion is that the event does affect security returns but that it does not contradict the efficient market hypothesis.

If, by contrast, the t -statistics continue to be significant on the post-event days, we might conclude that the market is inefficient in that it does not quickly absorb new information. We might also conclude that the market is inefficient if we were to observe significant t -statistics on the day of the event and we had reason to believe that the event (including its magnitude) was anticipated.

Issues in Measuring Events

When designing an event study, how to measure the event is not always clear. Suppose, for example, the event is an annual earnings announcement. The announcement that annual earnings are \$3.00 a share is meaningless unless this number is contrasted to the market's expectation about earnings. Moreover, the market's expectation will have been conditioned by earlier information releases pertaining to earnings. Therefore, the first issue in measuring the event is to disentangle the unanticipated component of the announcement from the expected component.

The unanticipated component of the event is likely to be positive for some securities and negative for others, and the test of significance may need to be conditioned on the direction of the event. This can be accomplished by partitioning the sample into a subsample of securities for which the event was positive and a subsample for which the event was negative.

Another issue with respect to the measurement of the event is the influence of confounding factors. Suppose the event is defined as the announcement of a change in dividend policy. For many securities, this announcement may coincide with an information release about earnings. This coincident information is called a confounding event—an event that might distort or camouflage the effect of the event of interest on the security's return.

Issues in Measuring Return

In my description of the steps involved in an event study, I isolated the security-specific component of return by using the market model. The returns must be normalized so that the expected value of their unanticipated component is equal to zero percent. It is perfectly acceptable that the expected value of the unanticipated component of return conditioned on the event not equal zero, and it is equally acceptable that the unanticipated component of return conditioned on the absence of the event be systematically nonzero. The probability-weighted sum of the unanticipated components of return must equal zero, however.

The market model is but one method for adjusting returns. Some event studies adjust returns by subtracting from them the average return of the securities during the pre-event period. This adjustment procedure is called the mean adjustment. An alternative procedure is to subtract the market's coincident return from the security's return. This adjustment procedure is called the market adjustment.

The procedure described earlier to normalize the unanticipated component of return to zero using the market model is called risk adjustment. Risk adjustment of returns can also be accomplished by using a procedure pioneered by Fama and MacBeth in 1973.² The unanticipated component of return is derived by computing an expected return in period t and then subtracting it from the security's actual return in period t .

The first step in this procedure is to estimate each security's beta by regressing its returns on the market's returns over some pre-event measurement period. Then, the returns across many securities in the same period t are regressed on their historical betas as of the beginning of period t . The intercept and slope from this cross-sectional regression are then used to measure the security's expected return.

Specifically, a security's expected return in period t is equal to the cross-sectional alpha in period t plus the cross-sectional beta in period t times the security's historical beta. The security's unanticipated component of return, therefore, equals its actual return in period t minus its expected return in period t (estimated from the cross-sectional coefficients and the security's historical beta).

The final approach for normalizing the unanticipated component of return to zero uses control portfolios. A control portfolio of sample securities is constructed to have a beta equal to 1. The

unanticipated component of return in an event-related period is computed as the return of the control portfolio less the return of the market.

Issues in Evaluating the Results

In the earlier example, a t -statistic was used to evaluate whether the event affected security returns. The use of a t -test presupposes that the returns of the securities from which the sample is drawn are normally distributed.

If we have reason to believe that the returns are not normally distributed, we can use a nonparametric test to evaluate the result. A nonparametric test, which is sometimes referred to as a distribution-free test, does not depend on the assumption of normality.

One of the simplest nonparametric tests is called a sign test. Not only is the sign test distribution free, it is also insensitive to the magnitude of the returns. It simply tests whether there are more positive returns (or negative returns, as the case may be) than would be expected if returns and the event are not related. This test statistic is computed as shown in Equation 6:

$$Z = \frac{(X - 0.5) - 0.5N}{0.5\sqrt{N}}, \quad (6)$$

where

Z = normal deviate

X = number of security-specific returns that are positive (or negative)

N = number of securities in sample

For example, if 13 returns are positive out of a sample of 20 securities, the normal deviate would equal 1.12, and we would fail to reject the null hypothesis that the event has no effect on security returns. If, instead, 65 returns are positive from a sample of 100 securities (which is the same proportion as 13 out of 20), the normal deviate would equal 2.90 and we would conclude that the event does affect security returns.

The sign test is but one of several nonparametric tests that can be used when the assumption of normality is in doubt or when the data are limited to ordinal values.

The t -statistic also assumes that the returns across the sample of securities are independent of one another. In many cases, security returns may not be mutually independent, even after they are risk adjusted. Securities may have other common sources of risk besides their exposure to the market. Perhaps the market-adjusted returns of securities within the same industry are correlated with

each other. This type of cross-correlation is particularly common in event studies of mergers when the propensity for mergers is an industry-related phenomenon. Sometimes, the problem of cross-correlation can be remedied by embellishing the risk-adjustment procedure to account for the portion of return that arises from industry affiliation or from exposure to some other source of common risk.

The Brown and Warner Study

In a classic article evaluating event study methodology, Brown and Warner simulated various risk-adjustment procedures to determine their efficacy.³ They first applied various methodologies to samples of securities that were contrived to have no abnormal returns in order to determine whether a particular methodology would reject the null hypothesis when it was true (a Type I error). Then, they artificially induced abnormal returns in samples to determine whether a particular methodology would fail to reject the null hypothesis when it was false (a Type II error). Finally, they compared the various methodologies based on their power to detect abnormal performance. The residual of a Type II error measures the power of a particular methodology.⁴

Brown and Warner concluded that none of the more elaborate procedures to isolate security-specific returns improved upon the simple market-model adjustment and that some of these procedures did not even improve upon the mean-adjustment procedure. Their message was that a researcher's time would be spent more productively by identifying and measuring the event rather than by devising elaborate procedures for controlling risk.

Footnotes

1. E. Fama, L. Fisher, M. Jensen, and R. Roll, "The Adjustment of Stock Prices to New Information," *International Economic Review*, vol. 10, no. 1 (February 1969):1-21.
2. E. Fama and J. MacBeth, "Risk, Return and Equilibrium: Empirical Tests," *Journal of Political Economy*, vol. 81, no. 3 (May/June 1973):607-36.
3. S. Brown and J. Warner, "Measuring Security Price Performance," *Journal of Financial Economics*, vol. 8 (September 1980):205-58.
4. For a review of hypothesis testing, see M. Kritzman, "What Practitioners Need to Know about Hypothesis Testing," *Financial Analysts Journal*, vol. 50, no. 4 (July/August 1994): 18-22.